



[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)

 Search: ☒ The ACM Digital Library ☐ The Guide



[Feedback](#) [Report a problem](#) [Satisfaction survey](#)

 Terms used **storage area network** and **replica\$** and **user**

Found 4,885 of 175,083

Sort results by


[Save results to a Binder](#)
[Try an Advanced Search](#)
[Try this search in The ACM Guide](#)

Display results


[Search Tips](#)
☐ Open results in a new window

Results 1 - 20 of 200

 Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

Best 200 shown

 Relevance scale ☐ ☐ ☐ ☐ ☐

### 1 [Brave new topics 3: advanced methods for medical image retrieval & applications:](#)



#### [Data grid for large-scale medical image archive and analysis](#)

H. K. Huang, Aifeng Zhang, Brent Liu, Zheng Zhou, Jorge Documet, Nelson King, L. W. C. Chan

 November 2005 **Proceedings of the 13th annual ACM international conference on Multimedia MULTIMEDIA '05**

Publisher: ACM Press

Full text available: pdf(2.03 MB)

 Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Storage and retrieval technology for large-scale medical image systems has matured significantly during the past ten years but many implementations still lack cost-effective backup and recovery solutions. As an example, a PACS (Picture Archiving and Communication system) in a general medical center requires about 40 Terabytes of storage capacity for seven years. Despite many healthcare centers are relying on PACS for 24/7 clinical operation, current PACS lacks affordable fault-tolerance storage ...

**Keywords:** PACS, bone age assessment of children, computational services, data grid, fault-tolerance archive, grid computing, image analysis, image data mining

### 2 [A formal model for reasoning about adaptive QoS-enabled middleware](#)



Nalini Venkatasubramanian, Carolyn Talcott, Gul A. Agha

 January 2004 **ACM Transactions on Software Engineering and Methodology (TOSEM)**, Volume 13 Issue 1

Publisher: ACM Press


Full text available: pdf(1.42 MB)

 Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Systems that provide distributed multimedia services are subject to constant evolution; customizable middleware is required to effectively manage this change. Middleware services for resource management execute concurrently with each other, and with application activities, and can, therefore, potentially interfere with each other. To ensure cost-effective QoS in distributed multimedia systems, safe composability of resource management services is essential. In this article, we present a meta-arc ...

**Keywords:** Middleware services, actors, meta-object models, multimedia, quality-of-service, reflection, theoretical foundations

### 3 [Scalable and fault-tolerant support for variable bit-rate data in the exedra streaming](#)

**server**

Stergios V. Anastasiadis, Kenneth C. Sevcik, Michael Stumm  
November 2005 **ACM Transactions on Storage (TOS)**, Volume 1 Issue 4

**Publisher:** ACM Press

Full text available: pdf(1.01 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

We describe the design and implementation of the Exedra continuous media server, and experimentally evaluate alternative resource management policies using a prototype system that we built. Exedra has been designed to provide scalable and efficient support for variable bit-rate media streams whose compression efficiency leads to reduced storage space and bandwidth requirements in comparison to constant bit-rate streams of equivalent quality. We examine alternative disk striping policies, and qua ...

**Keywords:** Content distribution, multimedia compression

4 Data-centric storage in sensornets with GHT, a geographic hash table

Sylvia Ratnasamy, Brad Karp, Scott Shenker, Deborah Estrin, Ramesh Govindan, Li Yin, Fang Yu

August 2003 **Mobile Networks and Applications**, Volume 8 Issue 4

**Publisher:** Kluwer Academic Publishers

Full text available: pdf(255.10 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Making effective use of the vast amounts of data gathered by large-scale sensor networks (sensornets) will require scalable, self-organizing, and energy-efficient data dissemination algorithms. For sensornets, where the content of the data is more important than the identity of the node that gathers them, researchers have found it useful to move away from the Internet's point-to-point communication abstraction and instead adopt abstractions that are more data-centric. This approach entails *na* ...

**Keywords:** *algorithms, distributed systems, performance, sensor networks*

5 An end-to-end approach to globally scalable network storage



Micah Beck, Terry Moore, James S. Plank

August 2002 **ACM SIGCOMM Computer Communication Review , Proceedings of the 2002 conference on Applications, technologies, architectures, and protocols for computer communications SIGCOMM '02**, Volume 32 Issue 4

**Publisher:** ACM Press

Full text available: pdf(286.82 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

This paper discusses the application of end-to-end design principles, which are characteristic of the architecture of the Internet, to network storage. While putting storage into the network fabric may seem to contradict end-to-end arguments, we try to show not only that there is no contradiction, but also that adherence to such an approach is the key to achieving true scalability of shared network storage. After discussing end-to-end arguments with respect to several properties of network stora ...

**Keywords:** IBP, asynchronous communications, end-to-end design, exNode, internet backplane protocol, logistical networking, network storage, scalability, store and forward network, wide area storage

6 An Efficient Data Location Protocol for Self.organizing Storage Clusters

Hong Tang, Tao Yang

November 2003 **Proceedings of the 2003 ACM/IEEE conference on Supercomputing**

**Publisher:** IEEE Computer Society

Full text available:  [pdf\(345.61 KB\)](#) Additional Information: [full citation](#), [abstract](#)

Component additions and failures are common for large-scale storage clusters in production environments. To improve availability and manageability, we investigate and compare data location schemes for a large self-organizing storage cluster that can quickly adapt to the additions or departures of storage nodes. We further present an efficient location scheme that differentiates between small and large file blocks for reduced management overhead compared to uniform strategies. In our protocol, sm ...

## 7 Interposed request routing for scalable network storage

 February 2002 **ACM Transactions on Computer Systems (TOCS)**, Volume 20 Issue 1

**Publisher:** ACM Press

Full text available:  [pdf\(363.12 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

This paper explores interposed request routing in Slice, a new storage system architecture for high-speed networks incorporating network-attached block storage. Slice interposes a request switching filter---called a *μproxy*---along each client's network path to the storage service (e.g., in a network adapter or switch). The *μproxy* intercepts request traffic and distributes it across a server ensemble. We propose request routing schemes for I/O and file service traffic, and explore th ...

**Keywords:** Content switch, file server, network file system, network storage, request redirection, service virtualization

## 8 FAB: building distributed enterprise disk arrays from commodity components

 Yasushi Saito, Svend Frølund, Alistair Veitch, Arif Merchant, Susan Spence

October 2004 **ACM SIGARCH Computer Architecture News , ACM SIGOPS Operating Systems Review , ACM SIGPLAN Notices , Proceedings of the 11th international conference on Architectural support for programming languages and operating systems ASPLOS-XI**, Volume 32 , 38 , 39 Issue 5 , 5 , 11


**Publisher:** ACM Press

Full text available:  [pdf\(671.67 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

This paper describes the design, implementation, and evaluation of a Federated Array of Bricks (FAB), a distributed disk array that provides the reliability of traditional enterprise arrays with lower cost and better scalability. FAB is built from a collection of *bricks*, small storage appliances containing commodity disks, CPU, NVRAM, and network interface cards. FAB deploys a new majority-voting-based algorithm to replicate or erasure-code logical blocks across bricks and a reconfigurati ...

**Keywords:** consensus, disk array, erasure coding, replication, storage, voting

## 9 Video Streaming 1: A Demand Adaptive and Locality Aware (DALA) streaming media server cluster architecture

 Zihui Ge, Ping Ji, Prashant Shenoy

May 2002 **Proceedings of the 12th international workshop on Network and operating systems support for digital audio and video**

**Publisher:** ACM Press

Full text available:  [pdf\(164.52 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The wide availability of broadband networking technologies such as cable modems and

DSL coupled with the growing popularity of the Internet has led to a dramatic increase in the availability and the use of online streaming media. With the "last mile" network bandwidth no longer a constraint, the bottleneck for video streaming has been pushed closer to the server. Streaming high quality audio and video to a myriad of clients imposes significant resource demands on the server. In this work, we pro ...

**Keywords:** demand adaptive, locality aware, server cluster, streaming media

10 DISP: Practical, efficient, secure and fault-tolerant distributed data storage



Daniel Ellard, James Megquier

February 2005 **ACM Transactions on Storage (TOS)**, Volume 1 Issue 1

**Publisher:** ACM Press

Full text available: pdf(148.11 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

DISP is a practical client-server protocol for the distributed storage of immutable data objects. Unlike most other contemporary protocols, DISP permits applications to make explicit tradeoffs between total storage space, computational overhead, and guarantees of availability, integrity, and privacy on a per-object basis. Applications specify the degree of redundancy with which each item is encoded, what level of integrity checks are computed and stored with each item, and whether items are stor ...

**Keywords:** Distributed data storage

11 Potpourri: Fast and transparent recovery for continuous availability of cluster-based servers



Rosalia Christodouloupoulou, Kaloian Manassiev, Angelos Bilas, Cristiana Amza

March 2006 **Proceedings of the eleventh ACM SIGPLAN symposium on Principles and practice of parallel programming PPOPP '06**

**Publisher:** ACM Press

Full text available: pdf(111.02 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Recently there has been renewed interest in building reliable servers that support continuous application operation. Besides maintaining system state consistent after a failure, one of the main challenges in achieving continuous operation is to provide fast reconfiguration. The complexity of the failure reconfiguration mechanisms employed and their overheads depend on the type of platform that is being used as a server and the types of applications that need to be supported. In this paper we foc ...

**Keywords:** availability, distributed shared memory, fast failure reconfiguration, fault tolerance, scalability

12 Per-user profile replication in mobile environments: algorithms, analysis, and simulation results



Narayanan Shivakumar, Jan Jannink, Jennifer Widom

October 1997 **Mobile Networks and Applications**, Volume 2 Issue 2

**Publisher:** Kluwer Academic Publishers

Full text available: pdf(1.84 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

We consider per-user profile replication as a mechanism for faster location lookup of mobile users in a personal communications service system. We present a minimum-cost maximum-flow based algorithm to compute the set of sites at which a user profile should be replicated given known calling and user mobility patterns. We show the costs and benefits of our replication algorithm against previous location lookup approaches through

analysis. We also simulate our algorithm against other location ...

### 13 Query optimization for vector space problems



K. Goda, T. Tamura, M. Kitsuregawa, A. Chowdhury, O. Frieder

September 2001 **Proceedings of the 24th annual international ACM SIGIR conference on Research and development in information retrieval**

**Publisher:** ACM Press

Full text available: pdf(83.57 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

We present performance measurement results for a parallel SQL based information retrieval system implemented on a PC cluster system. We used the Web-TREC dataset under a left-deep query execution plan. We achieved satisfactory speed up.

### 14 User profile replication for faster location lookup in mobile environments



Narayanan Shivakumar, Jennifer Widom

December 1995 **Proceedings of the 1st annual international conference on Mobile computing and networking**

**Publisher:** ACM Press

Full text available: pdf(915.70 KB) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

### 15 Partial replica selection based on relevance for information retrieval



Zhihong Lu, Kathryn S. McKinley

August 1999 **Proceedings of the 22nd annual international ACM SIGIR conference on Research and development in information retrieval**

**Publisher:** ACM Press

Full text available: pdf(271.29 KB) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

### 16 Data replicas in distributed information services



H. M. Gladney

March 1989 **ACM Transactions on Database Systems (TODS)**, Volume 14 Issue 1

**Publisher:** ACM Press

Full text available: pdf(1.94 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#), [review](#)

In an information distribution network in which records are repeatedly read, it is cost-effective to keep read-only copies in work locations. This paper presents a method of updating replicas that need not be immediately synchronized with the source data or with each other. The method allows an arbitrary mapping from source records to replica records. It is fail-safe, maximizes workstation autonomy, and is well suited to a network with slow, unreliable, and/or expensive communications links ...

### 17 REPLICAS - a new continuous system simulation language



Peter McLaughlin

March 1982 **Proceedings of the 15th annual symposium on Simulation**

Full text available: pdf(723.40 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

A new continuous system simulation language - REPLICAS, The Rational, Efficient Programming Language for the Implementation of Computerized Analysis and Simulation - is proposed for general engineering, scientific and econometric applications. The use of Gear's integration method coupled with a non-linear quasi-Newton solver relying on Broyden's method results in a reliable and efficient simulation system invoked by a language which requires only that the user define a mathematical model in ...

**18** Consistent and automatic replica regeneration

Haifeng Yu, Amin Vahdat

February 2005 **ACM Transactions on Storage (TOS)**, Volume 1 Issue 1**Publisher:** ACM PressFull text available:  pdf(372.24 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Reducing management costs and improving the availability of large-scale distributed systems require automatic replica *regeneration*, that is, creating new replicas in response to replica failures. A major challenge to regeneration is maintaining consistency when the replica group changes. Doing so is particularly difficult across the wide area where failure detection is complicated by network congestion and node overload. In this context, this article presents Om, the first read/write peer- ...

**Keywords:** Peer-to-peer storage systems, availability, consistency, regeneration, replication

**19** Replica allocation methods in ad hoc networks with data update

Takahiro Hara

August 2003 **Mobile Networks and Applications**, Volume 8 Issue 4**Publisher:** Kluwer Academic PublishersFull text available:  pdf(230.68 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

In ad hoc networks, since mobile hosts move freely, network division occurs frequently, and thus data accessibility is lower than that in conventional fixed networks. In this paper, assuming an environment where each data item is periodically updated, we propose three replica allocation methods to improve data accessibility by replicating data items on mobile hosts. In these three methods, we take into account the access frequency from mobile hosts to each data item, the status of the network co ...

**Keywords:** ad hoc networks, data accessibility, mobile computing environment, replica allocation

**20** The IceCube approach to the reconciliation of divergent replicas

Anne-Marie Kermarrec, Antony Rowstron, Marc Shapiro, Peter Druschel

August 2001 **Proceedings of the twentieth annual ACM symposium on Principles of distributed computing****Publisher:** ACM PressFull text available:  pdf(751.74 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

We describe a novel approach to log-based reconciliation called IceCube. It is general and is parameterised by application and object semantics. IceCube considers more flexible orderings and is designed to ease the burden of reconciliation on the application programmers. IceCube captures the static and dynamic reconciliation constraints between all pairs of actions, proposes schedules that satisfy the static constraints, and validates them against the dynamic constraints.

Preliminary ...

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2006 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)